

On quantum statistics in data analysis

Dusko Pavlovic*

Kestrel Institute and Oxford University
 dusko@{kestrel.edu, comlab.ox.ac.uk}

Abstract

Originally, quantum probability theory was developed to analyze statistical phenomena in quantum systems, where classical probability theory does not apply, because the lattice of measurable sets is not necessarily distributive. On the other hand, it is well known that the lattices of concepts, that arise in data analysis, are in general also non-distributive, albeit for completely different reasons. In his recent book, van Rijnsbergen (2004) argues that many of the logical tools developed for quantum systems are also suitable for applications in information retrieval. I explore the mathematical support for this idea on an abstract vector space model, covering several forms of data analysis (information retrieval, data mining, collaborative filtering, formal concept analysis...), and roughly based on an idea from categorical quantum mechanics (Abramsky & Coecke 2004; Coecke & Pavlovic 2007). It turns out that quantum (i.e., noncommutative) probability distributions arise already in this rudimentary mathematical framework. We show that a Bell-type inequality (Bell 1964) must be satisfied by the standard similarity measures, if they are used for preference predictions. The fact that already a very general, abstract version of the vector space model yields simple counterexamples for such inequalities seems to be an indicator of a genuine need for quantum statistics in data analysis.

Introduction

Until recently, Computer Science was mainly concerned with data storage and processing in purpose-built data bases and computers. With the advent of the Web and social computation, the task of finding and understanding information arising from local interactions in spontaneously evolving computational networks and data repositories has taken center stage.

As computers evolved from calculators, the key paradigm of Computer Science was computation-as-calculation, with the Turing Machine construed as a generic calculator, and with data processing performed by a small set of local operations. As computers got connected into networks, and captured a range of social functions, the paradigm of computation-as-communication emerged, with data processing performed not only locally, but also through distribution, merging, and association of data sets through vari-

ous communicating processes. Such non-local data processing has been implemented through markets, elections, and many other social mechanisms for a very long time, albeit on a smaller scale, with less concrete infrastructure, and with more complex computational agents. A new family of its implementations is based on a new computational platform, which is not any more the Computer, or even its operating system, but the Web, and its knowledge systems.

But while the interfaces of the local computational processes are defined to be the interfaces of the computers which perform them, the carriers of computation-as-communication do not come with clearly defined interfaces. The task of finding and supplying reliable data within a market, or on the Web, or in a social group, carries with it many deep problems. Two of them are particularly relevant for this work.

Problem of partial information and indeterminacy

Data processing in a network is ongoing. On the other hand, the data sets are usually incomplete, and information needs to be extracted from such incomplete sets. E.g., a task in a recommender system is to extrapolate which movies (books, music...) will a user like, from a sparse sample of those that she had previously rated. In information retrieval, the task is to extrapolate which information is relevant for a query, from a small set of tokens characterizing the query on one hand, and the information on the other hand.

In the standard model of data analysis, succinctly presented e.g. in (Azar *et al.* 2001), it is assumed that a matrix of random variables, containing a complete information about the relevant properties of the objects of interest, exists out there (in some sort of a Platonic heaven of information), and can be sampled. The problem of data analysis is that the sampling process is noisy, and partial; more specifically, that the distributions of the random variables are distorted by an error process, and by an omission process. The task of data analysis is to eliminate the effects of these processes, and reconstruct a good approximation of the original information.

While mathematically convenient, and computationally effective, this model does not seem very realistic. If we instantiate it to a recommender system again, then its basic assumption becomes that each user has a completely defined preference distribution, albeit only over the items that he has used, and that the recommender system just needs

*Supported by EPSRC and ONR.

to reconstruct this preference distribution. But if we zoom in, and ask the user himself, he will often be unable to precisely reconstruct his own preference distribution. If we ask him to rate some items again, he will often assign different ratings. One reason is that information processing is on-going, and that the preferences evolve and change. If we zoom in even further, we will find that the state of user's preferences is usually not completely determined even in a completely static model: right after watching a movie, one usually needs to toss a "mental coin" to decide whether to assign 2 or 3 stars, say, to the performance of an actor; or to decide whether to pay more attention, while watching the movie, to this or that aspect, music, colors...

While the indeterminacy of information in a network can be reduced to an effect of noise, like in the standard model, and averaged out, it is interesting to ponder whether viewing this indeterminacy as an essential feature of network computation, rather than a bug, may lead to more realistic models of information systems. Is the "mental coin", which resolves the superposition of the many components of my preferences when I need to measure them, akin to a real coin, which we all agree is governed by completely deterministic laws of classical physics, and its randomness is just the appearance of its complex behavior; or is this "mental coin" governed by a more fundamental form of randomness, like the one that occurs in quantum mechanics, causing the superposition of many states to *collapse under measurement*?

Problem of classification and latent semantics

The task of conceptualizing data has been formulated in many ways. In information retrieval, the central task is to determine the relevance of data with respect to a query. In recommender systems, the implicit query is always: "What will I like, given my past choices and rankings?", and the task is to find the relevant recommendations. In order to tackle such tasks, one classifies the data on one hand, the queries on the other, and aligns the two classifications, in order to extrapolate the future choices from the past choices. — But what are these classifications based on?

The simplest approach is based on keywords. But even classifying a corpus of purely textual documents, viewed as bags of words, according to the frequency of the occurrences of the relevant keywords, leads to significant problems: polysemy, homonymy, synonymy. The problem becomes very difficult when it comes to classifying families of non-textual objects: images, music, video, film. Only a small part of their correlations can be captured by connecting the keywords, captions, or other forms of textual annotations.

Latent semantics correlates data by extracting their intrinsic structure. For instance, the central piece of the original Google search engine, distinguishing it from other similar engines, was that the keyword search was supported by PageRank (Page *et al.* 1998), a reputation ranking of the Web pages, extracted from their intrinsic hyperlink structure. Even for the keyword search, the crucial step was to recognize this latent variable (Everitt 1984) extracting relevance from non-local network structure, rather than from local term occurrence. Such semantical support is even more critical for search and retrieval of non-textual information,

on the Web and in other data spaces.

Overview of latent semantics

We consider the case when two types of data assign the meaning to each other.

Pattern matrices

Latent semantics is generally given as a map

$$J \times U \xrightarrow{A} R$$

where

- J is a set of *objects*, or *items*,
- U is a set of *properties*, or *users*,
- R is a set of *values*, or *ratings*.

This map is conveniently presented as a *pattern matrix* $A = (A_{iu})_{J \times U}$. The entry A_{iu} can be intuitively written as a model relation $i \models u$, especially when $R = \{0, 1\}$. In general, it can be construed as the degree to which the object i satisfies the property, or the user u . While the ratings R usually carry a structure of an ordered *rig*¹, the attributes U often carry a more general algebraic structure, whereas the behaviors of the objects in J may be expressed coalgebraically. Clearly, the rig structure of R is just enough to support the usual matrix composition. Sometimes, but not always, we also assume that R has no nilpotents, so that it can be embedded in an ordered field.

Examples.

domain	J	U	R	A_{iu}
text analysis	documents	terms	\mathbb{N}	occurrence count
measurement	instances	quantities	\mathbb{R}	outcome
user preference	items	users	$\{0, \dots, 5\}$	rating
topic search	authorities	hubs	\mathbb{N}	hyperlinks
concept analysis	objects	attributes	$\{0, 1\}$	satisfaction
elections	candidates	voters	$\{0, \dots, n\}$	preference
market	producers	consumers	\mathbb{Z}	delivery
digital images	images	pixels	$[0, 1]$	intensity

Balancing and normalization

Notation. For every vector $x = (x_k)_{k=1}^n$, we define

- the average (expectation) $\mathbf{E}(x) = \frac{1}{n} \sum_{k=1}^n x_k$
- the ℓ_2 -norm $\|x\|_2 = \sqrt{\sum_{k=1}^n |x_k|^2}$,
- the ℓ_∞ -norm $\|x\|_\infty = \bigvee_{k=1}^n |x_k|$.

Item balancing of a semantics matrix A reduces each of its rows $A_{i\bullet}$, corresponding to the item i , to a row vector $A_{i\bullet}^0$, defined

$$A_{i\bullet}^0 = A_{i\bullet} - \mathbf{E}(A_{i\bullet})$$

The unassigned ratings in $A_{i\bullet}$ are padded by zeros.

In an item-balanced matrix records, the difference between the items with a higher average rating and the items

¹A *rig* $R = (R, +, \cdot, 0, 1)$ is a "ring without the negatives". This means that $(R, +, 0)$ and $(R, \cdot, 1)$ are commutative monoids satisfying $a(b + c) = ab + ac$ and $0a = 0$. The typical examples include natural numbers, non-negative reals, but also distributive lattices, which generally do not embed in a ring.

with a lower average rating is factored out. Only the satisfaction profile of each item is recorded, over the set of users who have assigned it better-than-average, or worse-than-average rating. The average and unassigned ratings are identified, and both become 0.

User balancing of a semantics matrix A reduces each of its columns $A_{\bullet u}$, corresponding to the user u , to a column vector $A_{\bullet u}^0$, with the expected value 0, by setting

$$A_{\bullet u}^0 = A_{\bullet u} - \mathbf{E}(A_{\bullet u})$$

The unassigned ratings are again padded by zeros.

In a user-balanced matrix, users' different rating habits, that some of them are more generous than others, are factored out. Only the satisfaction profile of each user is recorded, over the set of all items that she has rated. The average and unassigned ratings are identified, both with 0.

Item normalization of a semantics matrix A factors its rows into unit vectors; the **user normalization** factors its columns into unit vectors — by setting

$$\begin{aligned} \underline{A}_{i\bullet} &= \frac{A_{i\bullet}}{\|A_{i\bullet}\|_2} \\ \underline{A}_{\bullet u} &= \frac{A_{\bullet u}}{\|A_{\bullet u}\|_2} \end{aligned}$$

Comment. The purpose of balancing and normalization of raw semantic matrices is to factor out the aspects of rating that are irrelevant for the intended analysis. Whether a particular adjustment is appropriate or not depends on the intent, and on the available data. E.g., padding the available ratings by assigning the average rating to all unrated items may be useful in some cases, but it skews the data when the sample is small.² *In the rest of the paper, we assume that all such adjustments have been applied to data as appropriate, and we focus on the methods for extracting information from them.*

Classification


Through pattern matrices and latent semantics, the objects and the properties lend a meaning to each other. The simple method for extracting that meaning is based on the general ideas of Principal Component Analysis (Jolliffe 1986). This method underlies not only the vector space based approaches, like Latent Semantics Indexing (LSI) (Deerwester *et al.* 1990), or Hypertext Induced Topic Search (HITS) (Kleinberg 1999), but also, albeit in a less obvious way, Formal Concept Analysis (FCA) (Wille 1982), and some other approaches. The general idea is that the latent semantical structures can be obtained by factoring the pattern matrix through suitable transformations, required to preserve a *conceptual distance* between the objects, as well as between their properties. These distance-preserving transformations can be captured under the abstract notion of *isometry*.

²E.g., when only one rating is available from a user, then extrapolating his average rating to the unrated items simply erases all available information.


Suppose that the rig of values is given with an involutive automorphism $\overline{(-)} : R \rightarrow R$, called *conjugation*. If the values are the complex numbers, $R = \mathbb{C}$, then of course $\overline{a + ib} = a - ib$. For general rigs R , conjugation sometimes boils down to $\overline{a} = a$. In any case, any pattern matrix $A = (A_{iu})_{J \times U}$ induces an *adjoint* matrix $A^\dagger = (A_{ui}^\dagger)_{U \times J}$, whose entries are defined to be $A_{ui}^\dagger = \overline{A_{iu}}$. The *inner product* of vectors $x, y \in R^J$ can now be defined as $\langle x | y \rangle = y^\dagger \circ x$.

Definitions. An *isometry* is a map $U : \mathcal{A} \hookrightarrow \mathcal{B}$ such that $\langle Ux | Uy \rangle = \langle x | y \rangle$ holds for all x, y . Equivalently, this means that $U^\dagger U = \text{id}_{\mathcal{A}}$. It is a *unitary* if both U and U^\dagger are isometries.

An *isometric decomposition* of an operator $B : \mathcal{U} \rightarrow \mathcal{J}$ consists of isometries $\hat{V} : \hat{\mathcal{U}} \hookrightarrow \mathcal{U}$ and $\hat{W} : \hat{\mathcal{J}} \hookrightarrow \mathcal{J}$ such that there is a (necessarily unique) map $\hat{B} : \hat{\mathcal{J}} \rightarrow \hat{\mathcal{U}}$ satisfying $B = \hat{W} \hat{B} \hat{V}^\dagger$



The *spectral decomposition* $B = \bar{W} \bar{B} \bar{V}^\dagger$ is minimal among B 's isometric decompositions:



in the sense that for every isometric decomposition $B = \hat{W} \hat{B} \hat{V}^\dagger$, there is an isometric decomposition $\hat{B} = \bar{W} \bar{B} \bar{V}^\dagger$, such that $\bar{W} = \hat{W} \bar{W}$ and $\bar{V} = \hat{V} \bar{V}$.

We further also need

Correlation matrices are the self-adjoint matrices in the form $M^J = A A^\dagger$ and $M^U = A^\dagger A$, i.e.

$$\begin{aligned} M_{ij}^J &= \sum_{u \in U} \bar{A}_{ju} \cdot A_{iu} \\ M_{uv}^U &= \sum_{i \in J} A_{iu} \cdot \bar{A}_{iv} \end{aligned}$$

Examples of classification through isometric decomposition

Given a pattern matrix $J \times U \xrightarrow{A} R$, we set

$$\begin{aligned} \mathcal{J} &= R^J \\ \mathcal{U} &= R^U \end{aligned}$$

so that A becomes a linear operator $A : \mathcal{U} \rightarrow \mathcal{J}$, defined by the usual matrix action on the vectors.

Latent Semantic Indexing. (Deerwester *et al.* 1990) Let the rig of values \mathbb{R} be the field of real numbers \mathbb{R} , with the trivial conjugation $\bar{r} = r$. This means that $\mathcal{J} = \mathbb{R}^J$ and $\mathcal{U} = \mathbb{R}^U$ are real vector spaces. The pattern matrix $J \times U \xrightarrow{A} \mathbb{R}$ induces the linear operator $\mathcal{U} \xrightarrow{A} \mathcal{J}$ and the adjoint $\mathcal{J} \xrightarrow{A^\dagger} \mathcal{U}$ is just the transpose.

The isometric decomposition boils down to the singular value decomposition. The isometries $V : \mathcal{U}' \hookrightarrow \mathcal{U}$ and $W : \mathcal{J}' \hookrightarrow \mathcal{J}$ are obtained by the spectral decomposition of the symmetric matrices $M^U = A^\dagger A$ and $M^J = A A^\dagger$. Since both decompose through the same rank space, with the same spectrum $\Lambda = \{\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n\}$, we get a positive diagonal matrix Λ such that $A^\dagger A = V \Lambda V^\dagger$ and $A A^\dagger = W \Lambda W^\dagger$, from which $A = W D V^\dagger$ follows for $D = \sqrt{\Lambda}$.

The eigenspaces of M^U and M^J can be viewed as *pure topics* captured by the pattern matrix A . The eigenvalues correspond to the degree of semantical relevance of each topic in the data set from which the pattern matrix was extracted. If U are users and J items, then the eigenspaces in \mathcal{U} can be thought of as *tastes*, the eigenspaces in \mathcal{J} as *styles*. Remarkably, there is a bijective correspondence between the two, and the eigenvalues quantify the correlations. As an instance of the same decomposition, Kleinberg's (1999) analysis of Hyperlink Induced Topic Search (HITS) yields a similar correspondence between the hubs and the authorities on the Web. In all cases, the underlying view is that the information consumers and the information producers, lending each other the latent semantics, share a uniform conceptual space. An even simpler presentation of that optimistic view is

Formal Concept Analysis. (Ganter, Stumme, & Wille 2005) Let the rig of values \mathbb{R} now be the distributive lattice $\mathbb{B} = (2, \vee, \wedge, 0, 1)$, over the underlying set $2 = \{0, 1\}$, with the negation $\neg : \mathbb{B} \rightarrow \mathbb{B}$ as the conjugation $\bar{i} = \neg i$. Note that this is now an antimorphism of $\mathbb{B} = (2, \vee, \wedge, 0, 1)$ with the dual lattice $\tilde{\mathbb{B}} = (\tilde{2}, \wedge, \vee, 1, 0)$. The space of the objects is thus the boolean lattice $\mathcal{J} = 2^J$, ordered by inclusion, whereas the space of the properties is the boolean lattice $\mathcal{U} = \tilde{2}^U$, ordered by reverse inclusion.

Given a pattern matrix, which in this case boils down to a binary relation $J \times U \xrightarrow{A} 2$, we consider the induced map $\mathcal{U} \xrightarrow{A} 2^J$, and derive the monotone maps

$$B(X) = \{i \in J \mid \exists u \in X. \neg u A i\}$$

$$B^\dagger(Y) = \{u \in U \mid \forall i \notin Y. u A i\}$$

which are adjoint to each other in the sense

$$B(X) \subseteq Y \iff X \subseteq B^\dagger(Y)$$

and by conjugating yield the Galois connection

$$Y \subseteq \neg B(X) \iff X \subseteq B^\dagger(\neg Y)$$

$$\begin{array}{ccc} 2^U & \xrightleftharpoons[B^\dagger]{B} & 2^J \xrightleftharpoons[\neg]{\neg} 2^J \end{array}$$

The spectral decomposition

$$\begin{array}{ccc} 2^U & \xrightleftharpoons[B^\dagger]{B} & 2^J \end{array}$$

is obtained by setting

$$\bar{\mathcal{U}} = \{X \in 2^U \mid M^U(X) = X\}$$

$$\bar{\mathcal{J}} = \{Y \in 2^J \mid M^J(Y) = Y\}$$

where the closure operators $M^U = (B^\dagger \neg) \circ (\neg B)$ and $M^J = (\neg B) \circ (B^\dagger \neg)$ unfold to

$$M^U(X) = \{u \in U \mid \forall i \in J. (\forall v \in X. i A v) \Rightarrow i A u\}$$

$$M^J(Y) = \{i \in J \mid \forall u \in U. (\forall j \in Y. j A u) \Rightarrow i A u\}$$

Note that M^U is obtained by composing the matrices $\neg B$ and $B^\dagger \neg$ over the space $\tilde{2}^J$, where the composition \circ is dual to the usual one, i.e. $(P \circ Q)_{ik} = \bigwedge_j (P_{ij} \vee Q_{kj})$.

It is easy to see that the lattices of closed sets $\bar{\mathcal{U}}$ and $\bar{\mathcal{J}}$ are isomorphic, because they are both isomorphic with

$$\mathcal{L} = \{ \langle X, Y \rangle \in \mathcal{P}U \times \mathcal{P}J \mid B(X) = \neg Y \wedge B^\dagger(Y) = \neg X \}$$

This is the form in which a *concept lattice* is usually presented (Ganter, Stumme, & Wille 2005). The fact that the spectral composition is *minimal* means that it correlates users' *strongest tastes*, captured in $\bar{\mathcal{U}}$ with items' *strongest styles*, captured in $\bar{\mathcal{J}}$.

Remark. While LSI is a standard, well-studied data mining method, FCA has been less familiar in the data analysis communities, although an early proposal of a concept-lattice approach can be traced back to the earliest days of the information retrieval research (Salton 1968), predating both FCA and even the standard vector space model. More recently, though, the applications of FCA in information retrieval have been tested and explained (Carpineto & Romano 2004; Priss 2006; Poshyvanyk & Marcus 2007). The succinct presentation of LSI and FCA as special cases of the same pattern, in our abstract model above, points to the fact that the Singular Value Decomposition, on which LSI is based, and the Galois Connections, that lead to FCA, both subsume under the abstract structure of isometric decomposition, just instantiated to the rig of reals for LSI, and to the boolean rig for FCA. The simple structure of isometric decomposition, and the corresponding notion of conceptual distance, can thus be construed as the basic building block of semantical classification in data analysis. It turns out that already this rudimentary structure leads into quantum statistics.

Concept lattices are not distributive

While classical measures are defined over σ -algebras, which are distributive (and boolean) as lattices, quantum measures are defined over a more general family of algebras,

which need not be distributive lattices, but only orthomodular (Meyer 1986; Meyer 1993; Redei & Summers 2006).

A crucial, frequently made observation, eventually leading into quantum statistics, is that the lattices of concepts, and of topics, induced by the various forms of latent semantics, are *not distributive*. Indeed, since the lattice structure is induced by

$$\begin{aligned} x \wedge y &= x \cap y \\ x \vee y &= M(x \cup y) \end{aligned}$$

the closure operator M often disturbs the distributivity of the underlying set-theoretic operations. The observation that this non-distributivity of concept lattices lifts to the realm of information retrieval is due to van Rijsbergen. For reader's convenience, we repeat the intuitive example of $x \wedge (y \vee z) \neq (x \wedge y) \vee (x \wedge z)$ from (van Rijsbergen 2004, p. 36). In a taxonomy of animals, take $x = \text{"bird"}$, $y = \text{"human"}$ and $z = \text{"lizzard"}$. Then both $x \wedge y$ and $x \wedge z$ are empty, so that $(x \wedge y) \vee (x \wedge z)$ remains empty. On the other hand, $y \vee z = \text{"vertebrates"}$, because vertebrates are the smallest class including both humans and lizzards. Hence $x \wedge (y \vee z) = \text{"birds"}$ is not empty.

The point is that such phenomena arise from all forms of latent semantics. But beyond this point, there are even more specific indications of quantum statistics at work.

Similarity and ranking

At the core of the vector space model of information retrieval, data mining and other forms of data analysis lies the idea that the basic similarity measure, applicable to pairs of objects, or of attributes, or to the mixtures thereof, is expressible in terms of the inner product of their normalized (often also balanced) vectors:

$$\begin{aligned} s(i, j) &= \langle \underline{A}_{j\bullet} | \underline{A}_{i\bullet} \rangle = \sum_{u \in U} \underline{A}_{ju} \cdot \underline{A}_{iu} \\ s(u, v) &= \langle \underline{A}_{\bullet u} | \underline{A}_{\bullet v} \rangle = \sum_{i \in J} \underline{A}_{iu} \cdot \underline{A}_{iv} \end{aligned}$$

More generally, using the inner product one can also measure the similarity of pure topics x and y , viewed as linear combinations of the property vectors:

$$s_M(x, y) = \langle x | A^\dagger A | y \rangle = \langle Ax | Ay \rangle$$

In the same vein, the ranking of mixed topics, represented by the subspaces E of the space of properties, then corresponds to the trace operator:

$$\begin{aligned} \text{tr}_M(x) &= \langle x | A^\dagger A | x \rangle = \langle Ax | Ax \rangle \\ \text{tr}_M(E) &= \sum_{x \in B_E} \text{tr}_M(x) \end{aligned}$$

Noting that a correlation matrix $M = A^\dagger A$ amounts to what is in quantum statistics called an *observable*, we see that the ranking measures, already in the standard vector model, correspond to quantum measures. If the pattern matrices are furthermore normalized as to generate the correlation matrices with a unit trace, then they correspond to quantum probability distributions, or to quantum states.

Bell's inequality of similarities

In this final section, we attempt to use the described measure of similarity of users' tastes, derived from their past ratings of similar items, to predict the probability that they will agree in their future ratings. Although based on a simple, intuitive view of similarity and agreement, this prediction turns out to be impossible, as it leads to a contradiction. This impossibility result can be viewed as an indicator of a quantum statistical correlation, or at least as evidence that there is a problem with the straightforward statistical model of this simple situation.

The contradiction arises along the lines of Bell's derivation of his notable inequality (Bell 1964). More precisely, for any pair of users $x, y \in U$, represented by the unit vectors $x, y : J \rightarrow \mathbb{R}$, derive from their past ratings of the same items, we consider the random variables $X, Y : J' \rightarrow \{0, 1\}$, over a possibly larger set of items. Suppose that $X(i) = 1$ means that the user x likes the item i , and that $X(i) = 0$ means that she does not like it. We assume that the probability $P(X = Y) \in [0, 1]$ that X and Y will agree is proportional to their past similarity $s(x, y) \in [-1, 1]$, modulo the rescaling of $[-1, 1]$ to $[0, 1]$. This induces a constraint on the similarities.

Proposition. *Let the past preferences of $x_0, x_1, y_0, y_1 \in J$ be given as unit vectors $x_0, x_1, y_0, y_1 : U \rightarrow \mathbb{R}$. If the probability of their future agreement is determined by rescaling the similarity of their past preferences*

$$P(X = Y) = \frac{1 + s(x, y)}{2}$$

then their similarities must satisfy the following condition:

$$s(x_0, y_1) + s(x_1, y_1) + s(x_1, y_0) - s(x_0, y_0) \leq 2 \quad (1)$$

This follows from the general fact that the disagreement of $\{0, 1\}$ -valued random variables is a distance function.

Lemma. *Any three random variables $X, Y, Z : J \rightarrow \{0, 1\}$ satisfy*

$$P(X \neq Z) \leq P(X \neq Y) + P(Y \neq Z) \quad (2)$$

Proof. Let $W_{XY} : U \rightarrow \{0, 1\}$ be the random variable

$$W_{XY}(i) = \begin{cases} 1 & \text{if } X(i) \neq Y(i) \\ 0 & \text{if } X(i) = Y(i) \end{cases}$$

We claim that

$$W_{XZ} \leq W_{XY} + W_{YZ} \quad (3)$$

Towards the contradiction, suppose that there is $j \in J$ with $W_{XZ}(j) > W_{XY}(j) + W_{YZ}(j)$. This means that $W_{XZ}(j) = 1$, but $W_{XY}(j) = W_{YZ}(j) = 0$, and thus $X(j) \neq Z(j)$ but $X(j) = Y(j)$ and $Y(j) = Z(j)$ — which is clearly impossible. Therefore (3) must be true. But since $P(X \neq Y) = \mathbf{E}(W_{XY})$, averaging (3) gives (2). \square

Proof of the Proposition. Since $P(X = Y) = \frac{1 + s(x, y)}{2}$, it follows that $P(X \neq Y) = \frac{1 - s(x, y)}{2}$. Substituting this into (2) gives (1). \square

Corollary. *The probability of users' future agreement $P(X = Y)$ cannot be derived by rescaling the past similarities of their tastes $s(x, y)$, where the similarity measure s is defined by the inner product. The reason is that formula (1), which would have to be satisfied, does not always hold.*

Proof. The taste vectors $x_0 = (1, 0)$, $y_0 = (-1, 0)$, $x_1 = (-\frac{1}{2}, \frac{\sqrt{3}}{2})$ and $y_1 = (\frac{1}{2}, \frac{\sqrt{3}}{2})$ provide a counterexample for (1). \square

Interpretation. Why is it not justified to predict future agreements from past similarities, both defined in intuitively obvious ways? One line of explanation is that *the independence assumptions are violated*. As usually, the dependencies can be explained in terms of hidden variables (e.g., off-line interactions of the users), or in terms of non-local interactions. Another line of explanation is that *the dependencies are introduced in the model itself*. Intuitively, this means that the users, whose agreements are predicted, have not been sampled in the same measure space, and that their preferences should not be statistically mixed.

Remark. Rather than derived from similarity, users' semantical distance can be defined by $P(X = Y) = |Ax - Ay|_\infty$. A reader familiar with quantum probability theory (Meyer 1986; Meyer 1993) will recognize this interaction of the Hilbert space ℓ_2 and the Banach space ℓ_∞ , which acts on it as a von Neumann algebra, as the familiar interface between the quantum and the classical probabilities.

Conclusion and future work

We have shown that already in the basic, but sufficiently abstract models of information retrieval, data mining, and other forms of data analysis, a suitable version of Bell's argument applies, suggesting that the quantum statistical approach may be necessary.

The simple interpretation of Bell's argument is that the quantum statistical predictions refer to non-local interactions. More subtle interpretations lead into the issues of contextuality (Bell 1987, p. 9). In some cases, of course, both the non-local interactions and the contextual dependencies arise as a figment of the statistical model, mixing variables that cannot be sampled together. Either way, the version of the argument presented above suggests simple minded prediction based on the vector space model of information processing in a network may lead to problems if the locality of the interactions is not taken into account. Is it possible that genuine entanglement phenomena arise on a network?

After a moment of thought about this question, one gets a strange feeling that quantum probability might in fact be easier to comprehend in the realm of network computation, than in physics.³ While action at a distance is a highly unintuitive phenomenon in physics — Einstein called it "spooky" — in

³Perhaps like the theory of parallel universes, which seems to have more convincing interpretations in everyday life, and in distributed computation, than in physics.

network computation it can be reduced to the fact that the information may flow not only through the network links, but also off the network. This fact is not only intuitively natural, in the sense that, say, the data on the Web move not only in packets, along the Internet links, but they also get teleported from site to site, by people talking to each other, and then typing on their keyboards; but it is also information theoretically robust, in the sense that there are always covert channels. In abstract models, they can be represented in terms of non-local hidden variables, or in terms of entanglement. Either way, the operational content of quantum statistical methods will undoubtedly broaden the algorithmic horizons of network computation and data analysis, already by analyzing the meaning of the notable quantum algorithms in physics-free implementations. Convenient toolkits for combining quantum states, and for composing quantum operations (Coecke & Pavlovic 2007) are likely to acquire new roles in latent semantics. On the other hand, the generic no-cloning and no-broadcasting theorems (Barnum *et al.* 2006) are likely to point to some interesting statistical limitations, with a potential impact in security.⁴

Acknowledgement. I am grateful to Eleanor Rieffel for pointing out an error in an earlier version of this abstract, caused by some of my notational abuses.

References

- [Abramsky & Coecke 2004] Abramsky, S., and Coecke, B. 2004. A categorical semantics of quantum protocols. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science (LICS)*. IEEE Computer Society. Also arXiv:quant-ph/0402130.
- [Azar *et al.* 2001] Azar, Y.; Fiat, A.; Karlin, A. R.; McSherry, F.; and Saia, J. 2001. Spectral analysis of data. In *ACM Symposium on Theory of Computing*, 619–626.
- [Barnum *et al.* 2006] Barnum, H.; Barrett, J.; Leifer, M.; and Wilce, A. 2006. Cloning and broadcasting in generic probabilistic theories.
- [Bell 1964] Bell, J. S. 1964. On the Einstein-Podolsky-Rosen paradox. *Physics* 1:195–200.
- [Bell 1987] Bell, J. S. 1987. *Speakable and Unspeakable in Quantum Mechanics*. Cambridge University Press.
- [Carpineto & Romano 2004] Carpineto, C., and Romano, G. 2004. Exploiting the potential of concept lattices for information retrieval with credo. *Journal of Universal Computer Science* 10(8):985–1013.
- [Coecke & Pavlovic 2007] Coecke, B., and Pavlovic, D. 2007. Quantum measurements without sums. In Chen, G.; Kauffman, L.; and Lamnaco, S., eds., *Mathematics of Quantum Computing and Technology*. Taylor and Francis. 559–596.
- [Deerwester *et al.* 1990] Deerwester, S. C.; Dumais, S. T.; Landauer, T. K.; Furnas, G. W.; and Harshman, R. A. 1990.

⁴One direct consequence of the no-cloning theorem seems to be that only classical styles can be copied.

- Indexing by latent semantic analysis. *Journal of the American Society of Information Science* 41(6):391–407.
- [Everitt 1984] Everitt, B. 1984. *An Introduction to Latent Variable Models*. London: Chapman & Hall.
- [Ganter, Stumme, & Wille 2005] Ganter, B.; Stumme, G.; and Wille, R., eds. 2005. *Formal Concept Analysis, Foundations and Applications*, volume 3626 of *Lecture Notes in Computer Science*. Springer.
- [Jolliffe 1986] Jolliffe, I. T. 1986. *Principal Component Analysis*. Springer Series in Statistics. Springer-Verlag.
- [Kleinberg 1999] Kleinberg, J. M. 1999. Authoritative sources in a hyperlinked environment. *Journal of the ACM* 46(5):604–632.
- [Meyer 1986] Meyer, P.-A. 1986. Éléments de probabilités quantiques (exposés I à IV). In *Séminaire de probabilités de Strassbourg*, volume 1204,1247 of *Lecture Notes in Mathematics*. Berlin: Springer-Verlag.
- [Meyer 1993] Meyer, P.-A. 1993. *Quantum Probability for Probabilists*. Number 1538 in *Lecture Notes in Mathematics*. Springer-Verlag.
- [Page et al. 1998] Page, L.; Brin, S.; Motwani, R.; and Winograd, T. 1998. The PageRank citation ranking: Bringing order to the Web. Technical report, Stanford Digital Library Technologies Project.
- [Poshyvanyk & Marcus 2007] Poshyvanyk, D., and Marcus, A. 2007. Combining formal concept analysis with information retrieval for concept location in source code. In *ICPC '07: Proceedings of the 15th IEEE International Conference on Program Comprehension*, 37–48. Washington, DC, USA: IEEE Computer Society.
- [Priss 2006] Priss, U. 2006. Formal concept analysis in information science. In Cronin, B., ed., *Annual Review of Information Science and Technology*, volume 40.
- [Redei & Summers 2006] Redei, M., and Summers, S. J. 2006. Quantum probability theory. To appear in *Studies in the History and Philosophy of Modern Physics*.
- [Salton 1968] Salton, G. 1968. *Automatic Information Organization and Retrieval*. McGraw Hill Text.
- [van Rijsbergen 2004] van Rijsbergen, C. J. 2004. *The Geometry of Information Retrieval*. New York, NY, USA: Cambridge University Press.
- [Wille 1982] Wille, R. 1982. Restructuring lattice theory: an approach based on hierarchies of concepts. In Rival, I., ed., *Ordered Sets*. Dordrecht: Dan Reidel. 445–470.